

CLAIMS

WHAT IS CLAIMED IS:

1. A method for coordinating a computation upon a plurality of data containers deployed at a plurality of nodes, comprising:

at each of the nodes, executing a corresponding process configured for starting a program to perform the computation in response to a command received from a database system; and at the database system, receiving a statement specifying an external routine for performing the computation and, in response to receiving the statement:
concurrently transmitting a plurality of commands for performing the computation to each said corresponding process;
receiving results from each said process; and
completing processing of the statement based on the results received from each said process.

2. A method for coordinating a computation upon a plurality of data containers deployed at a plurality of nodes, comprising:

receiving a statement at a database system specifying an external routine for performing the computation; and
in response to receiving the statement:
transmitting a plurality of commands for performing the computation to a plurality of respective processes configured for starting a plurality of respective programs to perform the computation in response to the commands;
receiving results from each said process; and
completing processing of the statement based on the results received from each said process.

3. A method according to claim 2, further comprising:
determining a cohort of nodes from among the plurality of nodes capable of performing the computation,
wherein the plurality of the respective processes correspond to the cohort of the nodes.
4. A method according to claim 3, wherein the plurality of nodes includes at least one node not included in the cohort of the nodes.
5. A method according to claim 3, wherein said determining is based on a degree of parallelism supported by each of the nodes.
6. A method according to claim 3, further comprising accessing a registry specifying an association between the programs and the data containers,
wherein said determining is based on the association between the programs and the data containers.
7. A method according to claim 3, further comprising accessing a registry specifying respective attributes for the data containers,
wherein said determining is based on matching the respective attributes for the data containers with a parameter in the statement.
8. A method according to claim 3, further comprising accessing a registry specifying a partitioning function associated with the programs and the data containers,
wherein said determining is based on the results from executing the partitioning function associated with the programs and the data containers.

9. A method according to claim 2, wherein at least some of the programs execute in parallel.

10. A computer-readable medium bearing instructions for coordinating the computation, said instructions being arranged, upon execution by one or more processors to perform the method according to claim 1.